# A Survey on Reinforcement Learning in Aviation Applications

Pouria Razzaghi[1], Amin Tabrizian[1], Wei Guo[2], Shulu Chen[3], Abenezer Taye[1],
Ellis Thompson[2], Alexis Bregeon[4], Ali Baheri[5], Peng Wei[1]

## I. Introduction

### A. Preliminary

Compared with model-based control and optimization methods, reinforcement learning (RL) provides a data-driven, learning-based framework to formulate and solve sequential decision-making problems. The RL framework has become promising due to largely improved data availability and computing power in the aviation industry. Many aviation-based applications can be formulated or treated as sequential decision-making problems. Some of them are offline planning problems, while the others need to be solved in an online manner and are safety-critical. In this survey paper, we first describe standard RL formulations and solutions. Then we survey the landscape of existing RL-based applications in aviation. Finally, we summarize the paper, identify the technical gaps, and suggest future directions of RL research in aviation.

The RL methodology is comprehensively outlined in the remainder of this section. To begin, we briefly describe the RL problem formulation and a few key concepts. Following that, two classical categories of model-free RL algorithms will be presented: value-based and policy-based leanings. We then will present the more advanced techniques as well as modern actor-critic methods and multi-agent reinforcement learning (MARL). The overall structure of the RL methodology is shown in Fig. 1.

### B. Overview of Reinforcement Learning

Reinforcement learning is a branch of machine learning that is about an agent interacts with an environment to achieve a goal. The environment is stated in the form of a Markov decision process (MDP) used to solve sequential decision making problems. In an MDP problem, an agent takes a series of actions to maximize the total received reward from an unknown environment. This problem can be represented by a tuple of $(\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where $\mathcal{S}$ is the set of states, $\mathcal{A}$ is the set of actions, $P$ is the transition probability function $(P(s_{t+1}|s_t, a_t))$ that maps a state-action $(s_t, a_t)$ pair to a distribution of next possible states, $R$ is the received reward at each step, and $\gamma$ is the discount factor representing the relative importance of future and immediate rewards. The policy, $\pi(.)$, represents a mapping from an agent's state to a distribution on the action space. The optimal policy, $\pi^*(.)$, takes place where the summation of expected rewards, $\left(\sum_{i=0}^{\infty} \gamma^i r^{t+i}\right)$, over period of the course of action is maximized. Figure 2 depicts a block diagram of an RL process. An agent observes its current states and reward from

the environment; then, the agent selects an action according to its policy. This will change the state of the environment and the new reward and state in the next time step push back to the agent.

If the model, $P(s_t + 1|s_t, a_t)$, and reward $r_t$ are known, dynamic programming (DP) is one method for determining the optimal policy [1]. In the absence of a model, the agent should learn the optimal policy by observing the past interactions or by directly interacting with the environment, which is the RL problem (see details in Fig. 1). In recent decades, variations and improvements have been made to methods designed to solve real-world problems using RL formulations and solution algorithms.

One of the most important differentiation points in the RL algorithms is whether the agent has access to the model or not. Model-based RL algorithms learn and use the model, while those algorithms that do not consider the model are known as model-free. Although model-free methods do not benefit from the potential gains in sample efficiency that may come from using a model, they are also more straightforward to implement and tune. There are a few model-based RL algorithms such as world models [2], imagination-augmented agents (I2A) [3], model-based RL with model-free fine-tuning (MBMF) [4], model-based value expansion (MBVE) [5]. On the contrary, model-free methods have been extensively developed and more used recently. In the following, we describe some traditional and modern model-free RL methods.

### C. Standard Formulations of Model-Free Reinforcement Learning

*1) Value-based Methods:* Q-learning is one of the fundamental value-based RL algorithms introduced by Watkins [6] at the end of 1980s. A Q-value for every combination of state and action pair in an environment can be defined as Eq. 1. It represents an expected value of the cumulative reward at time step $t$ for an action $(a)$ when it follows a policy $\pi$ as follows:

$$Q_\pi(s, a) = \mathbb{E}_\pi\left[\sum_{i=0}^{\infty} \gamma^i r_{t+i+1}|s, a\right] \qquad (1)$$

where $i$ is the number of steps forward at time step $t$. After updating the Q-value, the algorithm attempts to determine how valuable it is to take a particular action in a specific state. A Q-table is made by all the stored Q-values of each state-action pair in a discrete space (a Q-function approximator is used for a continuous state space-action). The policy $\pi(s) = \arg\max Q(s, a)$ yields the highest total reward. An
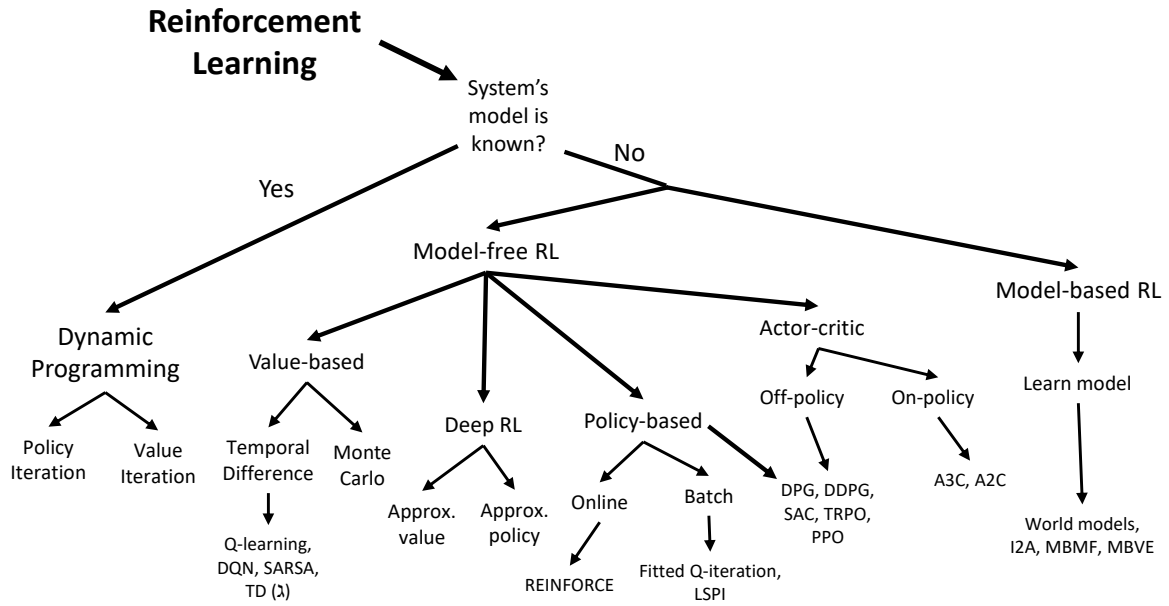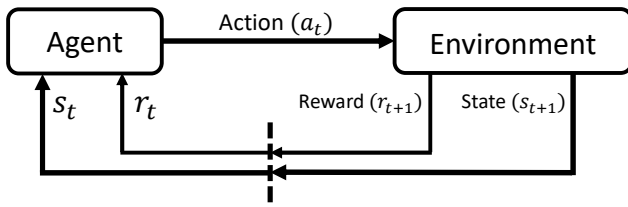
Fig. 1: Structure of the RL methodology.



Fig. 2: An agent selects an action $a_t$ based on its current state $s_t$, then it will receive a reward from the environment $r_t$ and arrive to the next state $s_{t+1}$. This process will go on till the agent arrives at a terminal state if any.

agent selects an action to explore the environment (so-called exploration-visiting almost all the state-action pairs a sufficient number of times) and observes the outcome. The Q-value can be updated by the temporal difference (TD) technique [7]:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \Big( r_t + \gamma \max_{a_{t+1} \in \mathcal{A}} Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t) \Big) \tag{2}$$

where $Q' = r_t + \gamma \max\limits_{a_{t+1} \in \mathcal{A}} Q(s_{t+1}, a_{t+1})$ is considered as the temporal difference target and $\alpha$ denotes the learning rate. We should note that Eq. (1) is a stochastic approximation scheme for the Bellman optimality equation solution and it will converge to $Q^*$ under certain assumptions [8], [9].

An off-policy method learns the value of the optimal policy independent of the agent's actions. Q-learning is considered an off-policy learning algorithm since it involves updating Q-value based on experiences that are not always generated from the derived policy. Whereas State–action–reward–state–action (SARSA) is an on-policy ones that generates experiences using the derived policy. For

example, SARSA uses $Q' = r_t + \gamma Q(s_{t+1}, a_{t+1})$ where $a_{t+1}$ is an action generated from the current policy or a given default policy.

A Monte Carlo method can be also used to estimate expected returns in non-Markovian episodic settings by averaging the results of multiple roll-outs. The Monte Carlo and TD methods have been joined and constructed the TD($\lambda$) [7]. The major problem of the traditional methods [10], [11] is the "curse of dimentionality". These methods rely on storing all the state-action pairs and representing in the tabular format that will grow exponentially as a factor of the number of states. One approach to solving this problem is to use a deep neural network (DNN) to approximate a parameterized Q-function. This creates a deep Q-networks (DQNs) [12]. DQN introduces replay memory and a separate target network, to overcome the problem of the instability and divergence issues in the training process of the approximation. To improve the stability of learning, this method uses a separate network $\hat{Q}$ for generating targets. A specific number of iterations is fixed for each episode. Moreover, by storing all transition experiences $(s_t, a_t, s_{t+1}, a_{t+1})$, the experience replay makes the randomly sampling for Q-learning updates more efficient [13]. In addition, the DQN performance is further improved by several notable variants, such as continuous DQN (cDQN) [14], double DQN [15], dueling DQN [16], and quantile regression DQN (QR-DQN) [17].

*2) Policy-based Methods:* Another family of RL algorithms are policy gradient algorithms, which do not calculate value but attempt to determine an optimal policy directly. In these algorithms, a probability distribution over a set of actions $(\pi(a|s, \theta))$ in relation to a policy defined as a function of parameters $\theta$ will be produced. An agent's likelihood of visiting state $s$ after applying a policy $\pi$ is described by the discounted state distribution. Using gradient ascent, the

policy is optimized for the objective function:

$$J(\theta) = \int_S \rho_\pi(s)\, r(s, \mu_\theta(s))\, ds = \mathbb{E}_{s_{\rho_\pi}}[r(s, \mu_\theta(s))] \quad (3)$$

where $\rho_\pi$ is the discounted state distribution [18]. The gradients are calculated $(\theta \leftarrow \theta + \eta \nabla J(\theta)$, where $\eta$ is the step size) while the actions are taken in accordance with the policy, and rewards are observed. More straightforwardly, the policy gradient methods choose actions directly from a model, then update the model weights to maximize the expected returns. The original policy-based method is called REINFORCE [19], which collects a full trajectory and then updates the policy weights in the Monte Carlo style and indicates that total return is sampled from the entire trajectory.

In the deterministic policy gradient (DPG) [20], instead of using a stochastic policy $(\pi(s, \theta))$, the actions are deterministically selected using policy $\mu(s, \theta)$. DPGs are limited cases of stochastic gradient policies when the variance becomes zero. The major drawback of a deterministic policy is the lack of exploration. For a proper exploration of the environment, the noise needs to be added and the policy becomes stochastic again $\big($adding Gaussian noise $\xi, a = \pi_\theta(s) + \xi\big)$. DPGs are therefore commonly implemented as actor-critic methods to allow off-policy exploration. Consequently, it is possible to add noise to action outputs for additional exploration without the need for a stochastic policy. Over action-value modeling, policy parametrization has the advantage of incorporating knowledge into the learning system in the form of the policy. Deep deterministic policy gradient (DDPG) [21] is a model-free off-policy algorithm for learning continuous actions, which combines ideas from DPG and DQN.

One of the problems of policy-based methods is in the gradient update. The policy performance drops if the updated policies deviate largely from previous ones. Trust region policy optimization (TRPO) [22] ensures a monotonic improvement in policy performance by optimizing a surrogate objective function. The policy gradient updates are enforced by approximating the Kullback-Leibler (KL) divergence between the old and new policies using a quadratic approximation to be in a given range. Proximal policy optimization (PPO) [23] achieves the same benefits as TRPO with a simplified implementation and improved sample complexity. It is revised based on TRPO, but only uses first-order optimization.

It is worth to mention that there is not a specific way to differentiate and easy-to-define the clusters of RL methods. Most of the aforementioned methods can be pointed out as an actor-critic architecture as it is illustrated in Fig. 1.

*3) Actor-critic Methods:* The actor-critic algorithm is an eminent and widely used architecture based on the combination of policy-based and value-based methods, inheriting their advantages [24]. The actor-critic algorithm can be considered as an TD learning method that represents the policy function independent of the value function. It introduces the eponymous components: the actor and the critic; the policy used to select actions is called the *actor*, and the estimated value function known as the *critic* criticizes the actions made by the actor [7].

The actor-critic methods achieved great success in many complex tasks; however, they suffer from various problems such as high variance, slow convergence, and local optimum. Therefore, many variants have been developed to improve the performance of actor-critic methods. Asynchronous advantage actor-critic (A3C) [25] uses advantage estimates rather than discounted returns in the actor-critic framework and asynchronously updates both the policy and value networks on multiple parallel threads of the environment. The parallel independent environments stabilize the learning process and enable more exploration. Advantage actor-critic (A2C) [26], the synchronous version of A3C, uses a single agent for simplicity or waits for each agent to finish its experience to collect multiple trajectories. This modification can significantly reduce the variance of the policy gradient estimate without changing the expectation. In this method, multiple actors are trained in parallel with different exploration policies, then the global parameters get updated based on all the learning results and synchronized to each actor. Soft actor-critic (SAC) [27] with stochastic policies is an off-policy deep actor-critic algorithm based on the maximum entropy RL framework. It benefits from adding an entropy term to the reward function to encourage better exploration.

*4) Multi-agent Reinforcement Learning Methods:* There are some new-born control tasks to regulate the behaviour of a multi-agent system interacting in a common environment. Multi-agent reinforcement learning (MARL) will be critical for the development of communication skills and other intellectual capacities, as well as teaching agents how to cooperate without causing harm to each other. These challenging tasks motivate researchers to use multi-agent RL frameworks [28]. A summary of related algorithms and theories is outlined in [29]. In the MARL framework, a set of $N$ agents interact with the same environment. At each time step and for a given state, each agent takes its own action, receiving a reward. The system then propagates to the next state. In MARL framework, multi-task and partially observation are usually considered [30]. The centralized and decentralized multi-agent RL methods attract much attention in aviation applications [31]. One popular variant involves each agent adopting a policy, which determines the action based on local observations. As only local observations are required for the execution, this method permits decentralized implementation. However, centralized training is still required since the system's state transition relies on the actions of every agent.

## II. Selected Applications of RL in Aviation

Many challenging problems in aviation can now be addressed using data-driven and machine-learning-based methods due to the availability of aviation data and significant increases in computational power. These problems are not limited to the following examples: air traffic management [32], aircraft sequencing [33], air traffic flow extraction [34], taxi-out time prediction [35], flight delay prediction [36],

[37], trajectory prediction [38], and aircraft performance parameter predicting [39].

The RL method as an area of machine learning have been a proper candidate to study the aviation problems. Figure 3 illustrates the taxonomy of RL in aviation. In the following sections, we try to summarize the usage of the RL in different applications. In the best of authors knowledge, this survey paper is the first study that review the RL methods in aviation.



Fig. 3: Taxonomy layout of RL in aviation.

## A. Collision Avoidance and Separation Assurance

Air traffic control (ATC) plays a crucial role in the air traffic management (ATM) system as it is responsible for maintaining flight safety and efficiency. Collision avoidance is the last layer of defense against mid-air collision. On one hand, air traffic controllers must maintain a safe separation distance between any two aircraft at all times. This function is called conflict resolution or separation assurance. On the other hand, an early adaptation of an in-air collision avoidance system was the Traffic Alert and Collision Avoidance System (TCAS) [51] and more recently Next-Generation Airborne Collision Avoidance System (ACAS-X) [52], [53]. The latter was built upon TCAS, introducing a partially observable Markov decision process (POMDP) for the problem formulation. It provides audible and visual warnings to pilots by evaluating the time to closest approach, to determine if a collision is likely to occur. Many studies have been recently conducted on RL based collision avoidance and separation assurance, which a selection is presented in Table I.

A MDP collision avoidance in a free flight airspace was introduced in [50]. In a 3D environment with both cooperative (aircraft actively trying to avoid others) and non-cooperative aircraft (those are not concerned with collision avoidance), the MDP formulation in a free flight was able to avoid collision between aircraft. In [49], a DRL method was implemented as an optimization to a collision avoidance problem.

Showing beyond human-level performance in many challenging problems, the collision avoidance problem of unmanned aerial vehicles (UAV) has been solved by implementing a DQN algorithm [40]. Deep Q-Learning from demonstrations (DQfD) and reward decomposition were implemented to provide interpretable aircraft collision avoidance solutions in [41]. A DQN technique was also applied for the collision avoidance of UAVs [40], change routes and speeds in NASA's Sector 33 [55], compute corrections on top of the existing collision avoidance approaches [51], [53], and unmanned free flight traffic in a dense airspace [49]. A framework using RL and GPS waypoints to avoid collisions was suggested in [56]. A double deep Q-network (DDQN) was applied to guide the aircraft through terminal sectors without collision in [57]. The approach tackles the cases where traditional collision avoidance methods fail namely in a dense airspace, those expected to be occupied by UAVs, and demonstrated the ability to provide reasonable corrections to maintain sufficient safety among aircraft systems.

The PPO methods are widely used in aircraft collision avoidance and have shown promising success. The problem of collision avoidance in structured airspace using PPO networks [42] was addressed using a Long Short-term Memory (LSTM) network [43], and attention networks [44] to handle a variable number of aircraft. While these algorithms show high performance in the training environment, a slight change in the evaluation environment can decrease the performance of these PPO models. A safety module based on Monte-Carlo Dropout [58] and execution-time data augmentation was proposed to solve the collision avoidance problem in environments, which are different from the training environments [45]. A PPO network was proposed for unmanned aircraft to provide a safe and efficient computational guidance of operations [59] and guided UAV in continuous state and action spaces to avoid collision with obstacles [46]. A message-passing network [60] was introduced to support collision avoidance.

A prior physical information of airplanes was injected to build a physics informed DRL algorithm for aircraft collision avoidance [61]. A reward engineering approach was proposed in [62] to support the PPO network to solve the collision avoidance problem in a 2D airspace.

Several studies have applied DDPG [21] to the aircraft collision avoidance problems. A DRL method was applied to resolve the conflict between two aircrafts with continuous action space in the presence of uncertainty based on DPG in [47]. Also, an intelligent interactive conflict solver was used to acquire ATCs' preferences and an RL agent to suggest conflict resolutions capturing those preferences [63]. Later, the DDPG algorithm dealt with air sectors with increased traffic [64]. A proper heading angle was obtained by the DDPG algorithm before the aircraft reached the boundary of the sector to avoid the collisions [65]. DDPG method was

TABLE I: Selection from literature on RL in collision avoidance. State/Action space (S/A Space) can be continuous (C), discrete (D), or mixed (M).

| Reference | S/A Space | Algorithm | Policy Class | Key Features |
|---|---|---|---|---|
| Wulfe [40] | D/D | Double DQN | $\epsilon$-greedy | Prioritized sampling, regularization, discretization of dynamics. |
| Hermans [41] | M/D | DQN | | Deep Q-learning from Demonstrations, Reward Decomposition. |
| Brittain and Wei [42]–[44] | M/D | PPO, Attention network, and LSTM | ANN | 1. Adopting a multi-agent framework to handle collision avoidance. 2. Using LSTM to enhance the performance of PPO. |
| Guo et al. [45] | M/D | PPO, Dropout | ANN | Using Monte-Carlo Dropout and data augmentation to improve the safety in unseen environments. |
| Hu et al. [46] | C/C | PPO | ANN | Developing continuous control for unmanned aircraft system. |
| Pham et al. [47] | C/C | DPG | ANN | Developing an air traffic scenario simulator. |
| Wang et al. [48] | C/C | K-Control Actor Critic | ANN | Two-dimensional continuous action selection. |
| Li *et al.* [49] | C/C | DPG | ANN | 1. Building on ACAS to provide corrections for dense airspace. 2. Handling dense airspace. |
| Bertram *et al.* [50] | C/C | PPO | ANN | 1. Introducing solution for high-density UAM airspace. 2. Using an MDP based trajectory planner to avoid cooperative and non-cooperative aircraft. 3. Adopting multi-agent system to handle large numbers of aircraft. |
| Herman [51] | D/D | DDQN | $\epsilon$-greedy | 1. Introducing an onboard collision avoidance tool for pilots. 2. Interrogating an airspace with rule-based logic. |
| Jeannin *et al.* [52] | M/D | DQN | - | Formal verification of ACAS-X. |
| Kochenderfer *et al.* [53] | M/D | PPO, Dropout | ANN | Building on TCAS using a numeric lookup optimized to a probabilistic model. |

also proposed to mitigate collisions in high-density scenarios and uncertainties in [66]. A mixed approach, which combines the traditional geometric resolution and the DDPG model, was proposed to avoid the conflicts [67]. Multi-agent deep deterministic policy gradient (MADDPG) was applied to pair-wisely solve the collisions between two aircrafts [68]. Another MADDPG-based conflict resolution method reduced the workloads of ATC and pilots in operation [69].

The actor-critic algorithms are also popular in this application. $K$-control actor-critic algorithm was proposed to detect conflict and resolution with a 2D continuous action space in [48]. A policy function returns a probability distribution over the actions that the agent can take based on the given state. A graph-based network for ATC in 3D unstructured airspace was built in [70] to manage the airspace by avoiding potential collisions and conflicts. A multi-layer RL model was proposed to guide an aircraft in a multi-dimensional goal problem [71]. Also, an LSTM network and an actor-critic model were used to avoid collisions for fixed-wing UAVs [72].

Besides the popular models, other RL methods were also implemented for collision avoidance. A message-passing-based decentralized computational guidance algorithm was proposed in [73], which used a multi-agent Monte Carlo tree search (MCTS) [74] formulation. It was also able to prevent loss of separation (LOS) for UAVs in an urban air mobility (UAM) setting. A highly efficient MDP-based decentralized algorithm was established to prevent conflict with cooperative and non-cooperative UAVs in the free flight airspace in [50]. The MuZero algorithm [75] was proposed to mitigate a collision in [76]. Difference rewards tool was applied in [77] and a graph convolutional reinforcement learning algorithm solved the multi-UAV conflict resolution problem [78].

Incorporating a DRL model to learn a collision avoidance strategy with training a NN simultaneously could reduce the learn-time and execute a more accurate model due to

removing the discretization problem [79]. Though DRL has shown great success in aircraft separation assurance, there are still a lot of unsolved problems. These problems create crucial obstacles to building DRL models in this safety-critical application in the real world. One major problem is validation. DRL models for aircraft separation have deep structures and complex input states. The complex architecture makes it difficult to verify the properties of DRL models using the traditional formal methods. Current work with formal methods can only validate very simple properties with shallow DRL models. The lack of validation limits the trustworthiness of these DRL models and their use in real-world applications.

Another important question is the gap between simulation and reality. DRL for aircraft separation assurance is trained with simulators because the real-world training is too expensive considering the potential damage. However, it is not possible to have a simulation mimic reality perfectly. The distribution shift between the simulation and reality may constrain the learning performance of the DRL models.

Besides these two issues, DRL for aircraft separation assurance also faces the problems of general DRL models. For example, DRL for separation assurance currently has a low sampling efficiency, which highly restricts the training speed. Also, the DRL for separation assurance model works as a black-box. It cannot provide explainable decision-making in this process.

### B. Air Traffic Flow Management

Traffic management is an encompassing term for any system that directly affects or is used to decide air traffic movements. The overarching aim of these systems is to reduce delay while maintaining operational safety of the airspace. Generally, air traffic flow and capacity management are part of a common air traffic service (ATS) and interface with either pilots directly or through ATC. Finally, all these designed systems can be considered through two classifica-

TABLE II: Selection from literature on RL in air traffic flow management.

| Reference | S/A Space | Algorithm | Policy Class | Key Features |
|-----------|-----------|-----------|--------------|--------------|
| Xi et al. [80] | C/C | DQN | ANN | 1. Flow management for UAM airspace. 2. Using a DQN with genetic algorithm to solve DCB problem. |
| Tang and Xu [81] | C/C | K-Control Actor Critic | ANN | 1. Rule-based time-step environment to mimic DCB process. 2. MARL framework to address credit assignment problem. |
| Spatharis et al. [82] | C/C | Hierarchical RL | $\epsilon$-greedy | 1. Hierarchical approach partitions task into hierarchies of states and actions. 2. Hierarchical methods improve on DCB in the pre-tactical stage. |

tions; designed systems for unmanned traffic management (UTM) and UAS operations, and designed for more conventional operations.

Air Traffic Flow and Management (ATFM) is a subset of traffic management that focuses on ensuring the available airspace capacity is used efficiently. The capacity can be influenced not only by the sector's size, shape, or altitude but also by stochastic variables like winds, weather, and emergencies or more constant variables like airport capacity and throughput. Demand capacity balancing (DCB) is a predictive method to ensure the efficient operation of airspace or ground operations. A collaborative approach was introduced to DCB utilizing: assigning delays, allowing alternative trajectories, using fixed airspace sectorization, or adjusting airspace sectorization to efficiently manage airspace [54]. Unlike other solutions, synchronized collaborative-demand capacity balancing (SC-DCB) seeks to relax the constraints of airspace configurations, with the outcome demonstrating a reduction in active sectors resulting in better utilization of the active ones. In recent work [80], RL techniques have been utilized to examine their efficiency in the UAM flow management, using a state space consisting of data retrieved from aircraft, weather, airspace capacity, and traffic density surveillance and training data constructed through a Post-Hoc system. Multi-agent approaches in flow management also emerged [81]–[83] to demonstrate that a MARL approach can be used in a dense traffic area (hotspots) successfully resolve these hotspots by taking either a holding, departure, or cooperation actions. The approach also results in an overall reduction in delay.

Ground delay programs (GDP) deal with an excessive number of flights reaching an airport serving as another air traffic flow management mechanism. Airports' ability to handle arrivals may be adversely affected by weather conditions. Issuing terminal traffic management initiatives (TMIs) is a technique for reducing the number of incoming aircraft to an airport for a short period. One type of this technique is ground delay program. A data-driven approach based on a multi-armed bandit framework was proposed for suggesting TMI actions [84]. This would be beneficial for human decision-makers to evaluate whether a suggested solution is reasonable or not. The suggestions were based on historical data of forecasted and observed demand and capacity, chosen TMI actions, and observed performance. The results showed that almost all proposed algorithms slightly outperform the historical actions. [85] proposed four methods for recommending strategic TMI parameters during

uncertain weather conditions. The first two methods were based on random exploration, while the others were using an $\epsilon$-greedy approach and a Softmax algorithm. The fast-time simulation results demonstrated the strong performance of the two latter methods relative to the others, and their potential to help with dealing with weather uncertainty. A comparison between behavioral cloning (BC) and inverse reinforcement learning (IRL) in predicting hourly expert GDP implementation actions was made in [86]. Historical data was used to predict GDP decisions on San Francisco and Newark international airports. The IRL method was proposed to reduce the complexity by only exploring the states in the data. The results demonstrated that BC is a more robust predictive performance than the IRL GDP-implemented models. The experiments also suggested that neither the BC nor the IRL models predict the relatively infrequent GDP initialization or cancellation events well, unlike Q-learning, which tends to provide accurate predicted times [87]. Better prediction of taxi-out times will improve taxiing management, which can benefit trajectory planning by using GDP to reduce congestion.

With airspace becoming denser due to higher traffic and the introduction of emerging UAS/UTM technologies, traffic management solutions will be needed to demonstrate the ability of adaptation to accommodate not only higher volumes and densities of air traffic but also to any new requirements imposed by this new classification of air traffic. Additionally, the safety and capacity of these systems will require formal verification and standardized validation, moving the field of RL in ATM away from the laboratory and being ready to be accepted by official bodies. Finally, there are still many unknowns about how the UTM/UAS airspace will be constructed, adding a further layer of complexity to solution design; new systems should entertain this notion and provide flexibility while the airspace is still being defined.

*C. Airline Revenue Management*

In 1970s, there was limited control over ticket pricing and network scheduling. If one airline company wanted to increase its fare, a permission from a federal agency called Civil Aeronautics Board (CAB) was needed. The pricing regulation at that time always led to a higher fare. Airline deregulation happened in 1979, which allowed companies to schedule and price freely. Consequently, airline revenue management (ARM) came out as a business practice to set prices when there is perishable inventory. The ARM is an airline company's strategy to maximize revenue by optimizing ticket prices and product availability. The classic

TABLE III: Selection from literature on RL in airline revenue management.

| Reference | Problem Type | Algorithm | Policy Class | Key Features |
|-----------|-------------|-----------|--------------|--------------|
| Gosavi *et al.* [88] | Single leg | Q-learning | $\epsilon$-greedy | Infinite time horizon under the average reward optimizing criterion. |
| Lawhead and Gosavi [89] | Single leg | Bounded Actor-Critic | $\epsilon$-greedy | Test two types of reward: discounted reward MDP and the average reward SMDP. |
| Bondoux *et al.* [90] | Single leg | DQN | $\epsilon$-greedy | Comparison between DQL and RMS |
| Shihab and Wei [91] | Single leg | DQN | $\epsilon$-greedy | Considering both cancellation and overbooking in the environment. |
| Wang *et al.* [92] | Single leg | DQN Actor-Critic | $\epsilon$-greedy | Combining quantity-based RM and price-based RM together. |
| Alamdari and Savard [93] | Single leg & Network | DQN | AGen | Greedily generate a set of "effective" actions to replace the original action space. |

ARM problem could be divided into two types, quantity-based and price-based revenue managements (RM) [94].

Quantity-based RM works on a predefined n-class fare structure and determines how many tickets are protected for each fare class. Also, it focuses on the capacity control of single and network flight legs. As a representative of the quantity-based RM, the expected marginal seat revenue (EMSR) models [95] are wildly used in the modern airline industry. The price-based RM focuses more on the dynamic pricing situation.

Traditional and widely used approaches for ARM systems are model-based and data-driven, which heavily depend on the accuracy of forecasting data such as passenger arrival distribution, willingness to pay (WTP), and cancellation rate. Recently, researchers have been considering applying the model-free learning-based methods on ARM, such as optimal control theory or RL. A research direction of using RL in ARM started in 2002 [88], where the $\lambda$-smart algorithm was designed to cast the single-leg ARM problem as a semi-Markov decision problem (SMDP) over an infinite time horizon under the average reward optimizing criterion. Later, a bounded actor-critic approach was applied on the same problem [89]. Both studies claimed that the model's performance was better than the EMSR model. A DRL model on ARM has been introduced to integrate the domain knowledge with a DNN trained on graphical processing units (GPUs) [90]. A DRL model was also applied to the inventory control problem, using DQN and considering both cancellation and overbooking in their environment [91]. Some other improvements to DRL models have also appeared in recent years. For example, an ARM problem was studied by combining quantity-based RM, and price-based RM [92], while the DRL was applied on both the single leg and network leg problems [93].

The previous learning-based approaches consider the game between passengers and airline companies. However, there is limited work regarding the competitive pricing process among different airline companies. We believe it will be an exciting topic with the development of multi-agent reinforcement learning.

### D. Aircraft Flight and Attitude Control

Attitude control of an aircraft can be challenging due to the system's nonlinearities, uncertainties, and noises acting upon the system, which are intrinsically present in the environment. Recently, researchers have aimed to develop advanced controllers based on RL algorithms. A selection of RL methods in attitude control applications is presented in Table IV.

These proposed controllers have been used in target tracking [96], [97], single/multi-agent obstacle avoidance [97], [98], vision-based landing [99], stabilization [100]–[103], visual servoing [104], and flat spin recovery [105].

In [103], it was shown that training a controller directly by RL, based on a nonlinear or unknown model, is feasible. The performance of the controllers based on different RL algorithms was also compared in [106]. The results showed that a DQN is more suitable for discrete tasks than policy gradient or DDPG, whereas DDPG was shown to perform better in more complex tasks. Also, a DQN method was used to design attitude control systems for aircraft [103], [106]. In addition, the DDPG-based controllers were established in [97], [106], [107], [110], [111]. An improved DDPG method was combined with transfer learning and a control system was developed to perform autonomous maneuvering target tracking [97]. A DDPG-based controller was also studied, guiding a UAV to a fixed position in a horizontal plane from any position, and attitude [110].

Other studies have been conducted using PPO methods [98], [101], [108]. An improved MARL algorithm was developed, named multi-agent joint proximal policy optimization (MAJPPO), to perform formation and obstacle avoidance. The controller has used a moving averaging method to make each agent obtain a centralized state value function [98]. By performing the experimental comparison, it was shown that the MAJPPO algorithm could better deal with partially observable environments. A PPO-based controller was designed for stabilizing a fixed-wing UAV [101]. The trained policy outperformed a PID controller regarding the number of iterations required for convergences. It was also shown that the RL controller could generalize severe environmental disturbances.

Since RL has achieved significant progress in attitude control, it has been considered a promising approach for designing optimal and robust controllers. However, there are still some challenges that should be addressed. The gap between simulations and natural environments was experimentally demonstrated [109], which required a new training approach. A controller learned to adapt to the difference between training models and real environments. Exploration

TABLE IV: Selection from literature on RL in attitude control.

| Reference | S/A Space | Algorithms | Policy Class | Key Features |
|---|---|---|---|---|
| Li *et al.* [96] | C/C | Actor-Critic | ANN | 1. Compensating for the actuator fault and system input saturation. 2. Proving system stability by Lyapunov theory. |
| Li *et al.* [97] | C/C | MMN-DDPG Transfer Learning | ANN | 1. Introducing exploratory noises and parameter-based transfer learning to improve speed and generalization. 2. Performing target tracking and obstacle avoidance precisely in uncertain environments. |
| Zhao *et al.* [98] | - | Multi-agent joint PPO (MAJPPO) | ANN | 1. Using a moving window averaging of state-valued function to deal with multi-agent coordination problems. 2. MAJPPO, a centralized training and distributed execution. |
| Lee *et al.* [99] | C/C | Actor-Critic | ANN | Using a simple PID controller for handling attitude and position of UAV and a DRL algorithm to generate proper commands. |
| Xian *et al.* [100] | C/C | Actor-Critic | ANN | 1. Compensating the error of actor-critic network by a robust nonlinear sliding mode control method. 2. Achieving a better control performance compared to LQR. |
| Zhen *et al.* [101] | - | PPO | ANN | Achieving more precise control comparing to a PID controller in the presence of disturbance. |
| Huang *et al.* [102] | C/C | Actor-Critic | ANN | Introducing an NN approximation to learn the optimal controller online with no information of model. |
| Huang *et al.* [103] | C/C | DDQN | ANN / $\epsilon$ − greedy | Proposing model can train the controller in time domain directly on nonlinear or unknown model. |
| Shi *et al.* [104] | D/D | Q-learning | TD | Taking Q-learning for adaptive servoing gain adjustment. |
| Kim *et al.* [105] | D/C | DQN | ANN | Covering both unusual attitude and stable spin mode recoveries. |
| Zuo *et al.* [106] | C/C | DQN, PG , DDPG | ANN / $\epsilon$-greedy | 1. Being more efficient and faster. 2. Handling continuous action space but not efficient enough. |
| Wang *et al.* [107] | C/C | DDPG | ANN | Using a normal distribution for having better exploration. |
| Bohn *et al.* [108] | C/C | PPO | ANN | 1. Converging faster than PID. 2. Generalizing to turbulent wind conditions. |
| Wada *et al.* [109] | C/C | Actor-Critic(A3C) LSTM | ANN | 1. Stability of the NNs in different delays. 2. Experimentally demonstrating the reality and simulation gap. |

and exploitation balance is another dilemma in RL. A normal distribution noise for exploring the environment was used at the start of the training process [107]. It also proposed using Uhlenbeck-Ornstein stochastic noise for future works.

### E. Fault Tolerant Controller

A fault is a change in a system's property or parameters that causes the system to behave differently from its design. In the other word, failure is a condition that prevents a system from functioning. A fault-tolerant controller (FTC) is a control strategy that aims to improve the performance of a system operating in degraded performance due to a fault [116]. FTCs are characterized as model-based or data-driven, based on the method used to develop the controllers. Model-based techniques necessitate knowledge of the system's model and parameters to design a fault-tolerant controller. On the contrary, data-driven approaches learn the FTC directly from system data. The fundamental problem of a model-based FTC approach is that its effectiveness depends on the system model's correctness, which is difficult to establish when system parameters can vary due to faults. Furthermore, complex systems necessitate complicated controllers, which, in turn, impacts the controllers' robustness. On the other hand, data-driven techniques utilize data to design FTC without knowing the system's dynamics. As a result, data-driven methods, particularly RL-basedz techniques, have recently gained a lot of attention.

Several approaches have been proposed in the literature to solve the FTC controller using RL. Different RL algorithms, including DDPG, TRPO, and PPO, have been used to develop FTC techniques for quadrotor attitude control [112].

The results indicated that among the developed RL-based fault-tolerant controllers, the trained PPO-based attitude controller outperformed a fully tuned PID controller in terms of rising time, peak velocities achieved, and total error among the trained set of controllers. A DPG-based technique with an integral compensator was adopted to develop a position-tracking controller for the quadrotor [113]. The approach employed a two-phased learning scheme, with a simplified model being utilized for offline learning and the learned policy being refined during flight. The results showed that the learned FTC is sufficiently robust to model errors and external disturbances. A DDPG-based fault-tolerant policy for position tracking of quadcopters was proposed in [114]. The framework operates so that it runs simultaneously with the model-based controller and only becomes active when the system's behavior changes from the normal operating condition.

One of the significant drawbacks of model-free RL-based FTC methods is that there is no guarantee of convergence. To overcome this problem, a model-based framework for position tracking of octocopters was proposed [115]. Four RL algorithms were proposed, PPO, DDPG, Twin-Delayed DDPG (TD3), and soft actor-critic (SAC). The results showed that PPO is more suitable for a fault-tolerant task.

### F. Aircraft flight planning

Flight and trajectory planning is a well-known aviation problem and is crucial. While airspace users want the most optimal trajectory to minimize a cost function, many constraints, such as ground obstacles, capacity limitations, or environmental threats, make this problem difficult to solve.

TABLE V: Selection from literature of RL on fault tolerant controller.

| Reference | Problem Type | S/A Space | Algorithm | Key Features |
|---|---|---|---|---|
| Koch *et al.* [112] | Attitude control | C/C | DDPG, TRPO,PPO | Training RL algorithms to perform end-to-end attitude control. |
| Wang *et al.* [113] | Position tracking | C/C | DPG | Integrating DPG with integral compensator and adopting a two-phased approach. |
| Fei *et al.* [114] | Position tracking | C/C | DDPG | Running simultaneously with the model-based controller. |
| Bhan *et al.* [115] | Position tracking | C/C | DDPG,TD3 SAC,PPO | 1. Estimating fault-related parameters using an estimator. 2. trainig several RL algorithms using the estimated parameters. |

TABLE VI: Selection from literature on RL in flight planning.

| Reference | Problem Type | Algorithms | Policy Class | Key Features |
|---|---|---|---|---|
| Ragi *et al.* [117] | UAV path planning | POMDP NBO approximation | - | 1. Dynamical environment. 2. Wind effects taken into account. |
| Zhang *et al.* [118] | UAV path planning | Geometric RL | - | Convergence of calculating the reward matrix theoretically proven. |
| Yan *et al.* [119] | UAV path planning | D3QN, DDQN, DQN | $\epsilon$-greedy | 1. Stage scenario for simulation. 2. DRL approach and comparison of methods. |
| Spatharis *et al.* [120], [121] | DCB problem | Independent RL Edged-based MARL Agent-based MARL | $\epsilon$-greedy | 1. Formulation as a Markov game. 2. Real-world scenarios. |
| Chen *et al.* [122] | DCB problem | DDQN, experience replay | Adaptive $\epsilon$-greedy | 1. Comparison with actual method used in operations 2. Decentralized Training with Decentralized Execution |
| Bertram *et al.* [123] | Flight plan scheduling | FastMDP | $\epsilon$-greedy | 1. Centralized or distributed flight plan scheduling. 2. Parallelization for large-scale scheduling. |

Several techniques, including rerouting or ground delay are proposed to mitigate traffic congestion in most cases. The ATM domain is essentially based on temporal operations, with a capacity supply and demand model to manage air traffic flows. This operation can lead to capacity imbalances and create hot spots in sectors when capacity (defined as the number of aircraft accepted in a given sector during a given period) is exceeded. The planning of a trajectory or flight of an aircraft can be done in several stages defined in the ATM domain; the strategic phase includes the planning of flights performed between one year and D-7, the pre-tactical phase takes place between D-7 and D-1, and finally, the tactical phase takes place on D-day. An RL planner shows to be a promising tool to solve the pre-flight planning problems in dangerous environments [124].

UAV's versatility in performing tasks ranging from terrain mapping to surveillance and military missions makes this problem a fundamental part of aircraft operations. One of many defined missions for UAVs is to fly over ground targets. The theory of POMDP was presented for military use, and nominal belief-state optimization (NBO) was used to find the optimal trajectory considering threats, wind effects, or other agents [117]. Also, an RL approach was proposed to use geometric information from the drone's environment and produce smoother and more feasible trajectories in a real time planning [118]. The dueling double deep Q-networks (D3QN), DDQN, and DQN methods have been compared in [119] to solve the path planning problem for an agent in the context of a dynamic environment where it faces an environmental threat.

An RL method was used to resolve these hot spots with traffic speed regulation [125]. An agent representing a fix (a 2D point in the sector) can regulate the flows. By improving the computational capabilities, flights have been considered as agents, and MARL methods [120] were proposed to solve these capacity problems. Various algorithms were also studied: independent learners, edge-MARL, and agent-based-MARL, based on Q-learning techniques. Hot spots were solved using GDP, in which flight departures are delayed, to shift the whole trajectory [121]. The results show that collaborative methods yield better results. In order to reduce the search space, a hierarchical MARL scheme was proposed to solve the demand-capacity balancing (DCB) problem with GDP [82], thus allowing the abstraction of time and state-action. Inspired by supervised learning, multiple supervised-MARL frameworks built on PPO were suggested [81], where the agents representing the flights have three actions: hold their departure, take-off, or cooperate. This study indicated that adding supervisors can help improve search and generalization abilities. DQN and decentralized training and decentralized execution (DTDE) combined with replay experience [122] were also used to solve the DCB problem. In addition, a multi-agent asynchronous advantage actor-critic (MAA3C) framework was constructed to resolve airspace hot spots within a proper ground delay [126].

All these works aim to reduce hot spots by delaying flights while minimizing average delays and ensuring good distribution. Still, they have not studied other trajectory planning techniques. An RL approach was proposed to select a low-level heuristic to mitigate the air traffic complexity [127]. Flight level allocation, staggered departure times, and en route path deviation reduced congestion. In a UAM concept, the pre-departure airspace reservation problem as an MDP was formulated [123]. The first-in-first-out (FIFO) principle and the fast-MDP algorithm provided a conflict-free trajectory at the strategic stage. The scheduler, allowing both

centralized and decentralized flight planning, takes advantage of the computing power and parallelization of GPUs to process a large number of flights. A Learning-to-Dispatch algorithm was proposed to maximize the air capacity under emergency situations such as hurricane disasters [128].

### G. Airline Maintenance

Maintenance scheduling is the process of planning when and what type of maintenance check should be performed on an aircraft. The maintenance tasks of airlines are usually grouped into four-letter checks (A, B, C, and D). The level of detail in the maintenance check of these groups is different. For example, A- and B-checks are considered light maintenance, and C- and D-check as heavy maintenance and more detailed inspection. Usually, weather conditions and flight disruptions cause deviation from the scheduled plan. These uncertainties make aircraft maintenance scheduling a challenging task.

A look-ahead approximate dynamic programming methodology was developed for aircraft maintenance check [129]. Its schedules minimized the wasted utilization interval between maintenance checks while reducing the need for additional maintenance slots. The methodology was tested with two case studies of maintenance data of an A320 family fleet. The developed method showed significant changes in scheduled maintenance times; it reduced the number of A-checks by 1.9%, the number of C-check by 9.8%, and the number of additional slots by 78.3% over four years.

An RL-based approach was proposed in [130] to solve the aircraft's long-term maintenance optimization problem. The proposed method uses information about the aircraft's future mission, repair cost, prognostics and health management, etc., to provide real-time, sequential maintenance decisions. The RL-driven approach outperforms three existing commonly used strategies in adjusting its decision principle based on the diverse data in several simulated maintenance scenarios. The integration of an RL model for Human–AI collaboration in maintenance planning and the visualization of the Condition-Based Maintenance indicators were proposed in [131]. Optimal maintenance decision-making in the presence of unexpected events was also developed.

### H. Safety and Certification of Reinforcement Learning

Safety is of utmost importance in safety-critical applications such as aviation systems. Recent promising results in RL have encouraged researchers to apply such techniques to many real-world applications. However, the certification of learning-based approaches, including RL in safety-critical applications, remains an open research question [132], [133]. Recent surveys provide a comprehensive overview of efforts toward safe RL for safety-critical applications [134]. While there has been a lot of research interest in safe RL, especially in the autonomous driving community [135]–[137], safe RL problem is still underexplored in the aviation research community. The application of safe RL in aviation systems has been studied from different angles. For instance, recently, a safe DRL approach was proposed for autonomous airborne collision avoidance systems [62]. From the conflict resolution perspective, soft actor-critic models were used during vertical maneuvers in a layered airspace [138]. In a similar line of research, a safe deep MARL framework can identify and resolve conflicts between aircraft in a high-density [42].

From the run-time assurance perspective, a run-time safety assurance approach casts the problem as an MDP framework and uses RL to solve it [139]. Similarly, the path planning problem was framed as MDP and utilized MCTS for safe and assured path planning [140]. To guarantee the safety of real-time autonomous flight operations, an MCTS algorithm was proposed along with Gaussian process regression and Bayesian optimization to discretize the continuous action space [141]. Furthermore, a reinforcement learning framework predicts and mitigates the potential loss of separation events in congested airspace [142]. Recently, a safety verification framework was presented for design-time and run-time assurance of learning-based components in aviation systems [133].

### III. Conclusion

In this paper, after a review of the most common RL techniques and the overall methodology and principles of them, a survey of the application of RL in aviation is proposed. Ranging from airline revenue management to aircraft altitude control, the use of RL methods has shown a great interest in the literature in the last decade. Indeed, with the increase of computational power and access to a large source of data, this data-driven approach has become widely studied. Whether it is collision avoidance, traffic management, or other aviation-related problems, these learning-based frameworks show promising results, and a variety of algorithms and techniques are often studied for a specific problem. The most advanced techniques such as DRL or DPG are used to deal with critical systems such as collision avoidance or to handle the increase of growing air traffic in traffic management and flight planning. However, differences between the simulated environment and real-world application or its black-box scheme can still be a hindrance to implementation in the aviation industry, constrained by numerous safety measures. The certification of such methods is then a crucial point for these innovative and disruptive applications in aviation and should be one of the focuses of research in this area.

### IV. Acknowledgment

## References

[1] D. Bertsekas, *Dynamic programming and optimal control: Volume I.* Athena scientific, 2012, vol. 1.

[2] D. Ha and J. Schmidhuber, "World models," *arXiv preprint arXiv:1803.10122*, 2018.

[3] S. Racanière, T. Weber, D. Reichert, L. Buesing, A. Guez, D. Jimenez Rezende, A. Puigdomènech Badia, O. Vinyals, N. Heess, Y. Li, *et al.*, "Imagination-augmented agents for deep reinforcement learning," *Advances in neural information processing systems*, vol. 30, 2017.

[4] A. Nagabandi, G. Kahn, R. S. Fearing, and S. Levine, "Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning," *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7559–7566, 2018.

[5] V. Feinberg, A. Wan, I. Stoica, M. I. Jordan, J. E. Gonzalez, and S. Levine, "Model-based value estimation for efficient model-free reinforcement learning," *arXiv preprint arXiv:1803.00101*, 2018.

[6] C. J. C. H. Watkins, "Learning from delayed rewards," Ph.D. dissertation, King's College, 1989.

[7] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction.* MIT press, 2018.

[8] J. N. Tsitsiklis, "Asynchronous stochastic approximation and q-learning," *Machine learning*, vol. 16, no. 3, pp. 185–202, 1994.

[9] G. Qu and A. Wierman, "Finite-time analysis of asynchronous stochastic approximation and q-learning," *Conference on Learning Theory*, pp. 3185–3205, 2020.

[10] C. J. Watkins and P. Dayan, "Q-learning," *Machine learning*, vol. 8, no. 3, pp. 279–292, 1992.

[11] G. A. Rummery and M. Niranjan, *On-line Q-learning using connectionist systems.* Citeseer, 1994, vol. 37.

[12] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.

[13] L.-J. Lin, "Self-improving reactive agents based on reinforcement learning, planning and teaching," *Machine learning*, vol. 8, no. 3, pp. 293–321, 1992.

[14] S. Gu, T. Lillicrap, I. Sutskever, and S. Levine, "Continuous deep q-learning with model-based acceleration," *International conference on machine learning*, pp. 2829–2838, 2016.

[15] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double q-learning," *Proceedings of the AAAI conference on artificial intelligence*, vol. 30, no. 1, 2016.

[16] Z. Wang, T. Schaul, M. Hessel, H. Hasselt, M. Lanctot, and N. Freitas, "Dueling network architectures for deep reinforcement learning," *International conference on machine learning*, pp. 1995–2003, 2016.

[17] J. Duan, Y. Guan, S. E. Li, Y. Ren, Q. Sun, and B. Cheng, "Distributional soft actor-critic: Off-policy reinforcement learning for addressing value estimation errors," *IEEE Transactions on Neural Networks and Learning Systems*, 2021.

[18] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," *Advances in neural information processing systems*, vol. 12, 1999.

[19] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine learning*, vol. 8, no. 3, pp. 229–256, 1992.

[20] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller, "Deterministic policy gradient algorithms," *International conference on machine learning*, pp. 387–395, 2014.

[21] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *ICRA (Poster)*, 2016.

[22] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," *International conference on machine learning*, pp. 1889–1897, 2015.

[23] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[24] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "Lightgbm: A highly efficient gradient boosting decision tree," *Advances in neural information processing systems*, vol. 30, pp. 3146–3154, 2017.

[25] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," *International conference on machine learning*, pp. 1928–1937, 2016.

[26] J. X. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J. Z. Leibo, R. Munos, C. Blundell, D. Kumaran, and M. Botvinick, "Learning to reinforcement learn," *arXiv preprint arXiv:1611.05763*, 2016.

[27] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," *International conference on machine learning*, pp. 1861–1870, 2018.

[28] S. Gronauer and K. Diepold, "Multi-agent deep reinforcement learning: a survey," *Artificial Intelligence Review*, vol. 55, no. 2, pp. 895–943, 2022.

[29] K. Zhang, Z. Yang, and T. Başar, "Multi-agent reinforcement learning: A selective overview of theories and algorithms," *Handbook of Reinforcement Learning and Control*, pp. 321–384, 2021.

[30] S. Omidshafiei, J. Pazis, C. Amato, J. P. How, and J. Vian, "Deep decentralized multi-task multi-agent reinforcement learning under partial observability," *International Conference on Machine Learning*, pp. 2681–2690, 2017.

[31] W. Wang, Y. Liu, R. Srikant, and L. Ying, "3m-rl: Multi-resolution, multi-agent, mean-field reinforcement learning for autonomous uav routing," *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[32] M. Schmidt, "A review of aircraft turnaround operations and simulations," *Progress in Aerospace Sciences*, vol. 92, pp. 25–38, 2017.

[33] M. S. Ahmed, S. Alam, and M. Barlow, "A cooperative co-evolutionary optimisation model for best-fit aircraft sequence and feasible runway configuration in a multi-runway airport," *Aerospace*, vol. 5, no. 3, p. 85, 2018.

[34] M. Conde Rocha Murca, R. DeLaura, R. J. Hansman, R. Jordan, T. Reynolds, and H. Balakrishnan, "Trajectory clustering and classification for characterization of air traffic flows," *16th AIAA Aviation Technology, Integration, and Operations Conference*, p. 3760, 2016.

[35] H. Lee, W. Malik, and Y. C. Jung, "Taxi-out time prediction for departures at charlotte airport using machine learning techniques," *16th AIAA Aviation Technology, Integration, and Operations Conference*, p. 3910, 2016.

[36] N. Takeichi, R. Kaida, A. Shimomura, and T. Yamauchi, "Prediction of delay due to air traffic control by machine learning," *AIAA modeling and simulation technologies conference*, p. 1323, 2017.

[37] S. Choi, Y. J. Kim, S. Briceno, and D. Mavris, "Prediction of weather-induced airline delays based on machine learning algorithms," *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*, pp. 1–6, 2016.

[38] S. Ayhan and H. Samet, "Aircraft trajectory prediction made easy with predictive analytics," *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 21–30, 2016.

[39] R. Alligier, D. Gianazza, and N. Durand, "Machine learning and mass estimation methods for ground-based aircraft climb prediction," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 6, pp. 3138–3149, 2015.

[40] B. Wulfe, "Uav collision avoidance policy optimization with deep reinforcement learning," Stanford University, Tech. Rep., 2017.

[41] M. Herman, "Towards explainable automation for air traffic control using deep q-learning from demonstrations and reward decomposition," Master's thesis, Aerospace Engineering, 2021.

[42] M. Brittain and P. Wei, "Autonomous separation assurance in an high-density en route sector: A deep multi-agent reinforcement learning approach," *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 3256–3262, 2019.

[43] M. Brittain, X. Yang, and P. Wei, "A deep multi-agent reinforcement learning approach to autonomous separation assurance," *AIAA Journal of Aerospace Information Systems*, vol. 18, no. 12, 2021.

[44] M. W. Brittain and P. Wei, "One to any: Distributed conflict resolution with deep multi-agent reinforcement learning and long short-term memory," *AIAA Scitech 2021 Forum*, p. 1952, 2021.

[45] W. Guo, M. Brittain, and P. Wei, "Safety enhancement for deep reinforcement learning in autonomous separation assurance," *arXiv preprint arXiv:2105.02331*, 2021.

[46] J. Hu, X. Yang, W. Wang, P. Wei, L. Ying, and Y. Liu, "Obstacle avoidance for uas in continuous action space using deep reinforcement learning," *IEEE Access*, vol. 10, pp. 90 623–90 634, 2022.

[47] D.-T. Pham, N. P. Tran, S. K. Goh, S. Alam, and V. Duong, "Reinforcement learning for two-aircraft conflict resolution in the presence of uncertainty," *2019 IEEE-RIVF International Conference on Computing and Communication Technologies (RIVF)*, pp. 1–6, 2019.

[48] Z. Wang, H. Li, J. Wang, and F. Shen, "Deep reinforcement learning based conflict detection and resolution in air traffic control," *IET Intelligent Transport Systems*, vol. 13, no. 6, pp. 1041–1047, 2019.

[49] S. Li, M. Egorov, and M. Kochenderfer, "Optimizing collision avoidance in dense airspace using deep reinforcement learning," *Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM)*, 2019.

[50] J. Bertram and P. Wei, "Distributed computational guidance for high-density urban air mobility with cooperative and non-cooperative collision avoidance," *AIAA Scitech 2020 Forum*, p. 1371, 2020.

[51] W. H. Harman, "Tcas- a system for preventing midair collisions," *The Lincoln Laboratory Journal*, vol. 2, no. 3, pp. 437–457, 1989.

[52] J.-B. Jeannin, K. Ghorbal, Y. Kouskoulas, R. Gardner, A. Schmidt, E. Zawadzki, and A. Platzer, "Formal verification of acas x, an industrial airborne collision avoidance system," *2015 International Conference on Embedded Software (EMSOFT)*, pp. 127–136, 2015.

[53] M. J. Kochenderfer, J. E. Holland, and J. P. Chryssanthacopoulos, "Next-generation airborne collision avoidance system," Massachusetts Institute of Technology-Lincoln Laboratory Lexington United States, Tech. Rep., 2012.

[54] Y. Xu, X. Prats, and D. Delahaye, "Synchronised demand-capacity balancing in collaborative air traffic flow management," *Transportation Research Part C: Emerging Technologies*, vol. 114, pp. 359–376, 2020.

[55] M. W. Brittain and P. Wei, "Towards autonomous air trac control for sequencing and separation-a deep reinforcement learning approach," *2018 Aviation Technology, Integration, and Operations Conference*, p. 3664, 2018.

[56] B. Jacob, A. Kaushik, P. Velavan, and M. Sharma, "Autonomous drones for medical assistance using reinforcement learning," *Advances in Augmented Reality and Virtual Reality*, pp. 133–156, 2022.

[57] Q. Xu, J. Huang, Z. Liu, and H. Ding, "A method based on deep reinforcement learning to generate control strategy for aircrafts in terminal sector," *Artificial Intelligence in China*, pp. 356–363, 2021.

[58] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," *international conference on machine learning*, pp. 1050–1059, 2016.

[59] J. Hu and Y. Liu, "Uas conflict resolution integrating a risk-based operational safety bound as airspace reservation with reinforcement learning," *AIAA Scitech 2020 Forum*, p. 1372, 2020.

[60] R. Dalmau and E. Allard, "Air traffic control using message passing neural networks and multi-agent reinforcement learning," *Proceedings of the 10th SESAR Innovation Days, Virtual Event*, pp. 7–10, 2020.

[61] P. Zhao and Y. Liu, "Physics informed deep reinforcement learning for aircraft conflict resolution," *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[62] C. Panoutsakopoulos, B. Yuksek, G. Inalhan, and A. Tsourdos, "Towards safe deep reinforcement learning for autonomous airborne collision avoidance systems," *AIAA SCITECH 2022 Forum*, p. 2102, 2022.

[63] N. P. Tran, D.-T. Pham, S. K. Goh, S. Alam, and V. Duong, "An intelligent interactive conflict solver incorporating air traffic controllers' preferences using reinforcement learning," *2019 Integrated Communications, Navigation and Surveillance Conference (ICNS)*, pp. 1–8, 2019.

[64] P. N. Tran, D.-T. Pham, S. K. Goh, S. Alam, and V. Duong, "An interactive conflict solver for learning air traffic conflict resolutions," *Journal of Aerospace Information Systems*, vol. 17, no. 6, pp. 271–277, 2020.

[65] H. Wen, H. Li, Z. Wang, X. Hou, and K. He, "Application of ddpg-based collision avoidance algorithm in air traffic control," *2019 12th International Symposium on Computational Intelligence and Design (ISCID)*, vol. 1, pp. 130–133, 2019.

[66] D.-T. Pham, N. P. Tran, S. Alam, V. Duong, and D. Delahaye, "A machine learning approach for conflict resolution in dense traffic scenarios with uncertainties," *ATM Seminar 2019, 13th USA/Europe ATM R&D Seminar*, 2019.

[67] M. Ribeiro, J. Ellerbroek, and J. Hoekstra, "Determining optimal conflict avoidance manoeuvres at high densities with reinforcement learning," *Proceedings of the Tenth SESAR Innovation Days, Virtual Conference*, pp. 7–10, 2020.

[68] R. Isufaj, D. Aranega Sebastia, and M. A. Piera, "Towards conflict resolution with deep multi-agent reinforcement learning," *Proceedings of the 14th USA/Europe Air Traffic Management Research and Development Seminar (ATM2021), New Orleans, LA, USA*, pp. 20–24, 2021.

[69] J. Lai, K. Cai, Z. Liu, and Y. Yang, "A multi-agent reinforcement learning approach for conflict resolution in dense traffic scenarios," *2021 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC)*, pp. 1–9, 2021.

[70] J. Mollinga and H. van Hoof, "An autonomous free airspace en-route controller using deep reinforcement learning techniques," *ICRAT*, 2020.

[71] W. Zu, H. Yang, R. Liu, and Y. Ji, "A multi-dimensional goal aircraft guidance approach based on reinforcement learning with a reward shaping algorithm," *Sensors*, vol. 21, no. 16, p. 5643, 2021.

[72] Y. Zhao, J. Guo, C. Bai, and H. Zheng, "Reinforcement learning-based collision avoidance guidance algorithm for fixed-wing uavs," *Complexity*, vol. 2021, 2021.

[73] X. Yang and P. Wei, "Scalable multi-agent computational guidance with separation assurance for autonomous urban air mobility," *Journal of Guidance, Control, and Dynamics*, vol. 43, no. 8, pp. 1473–1486, 2020.

[74] G. Chaslot, S. Bakkes, I. Szita, and P. Spronck, "Monte-carlo tree search: A new framework for game ai." *AIIDE*, vol. 8, pp. 216–217, 2008.

[75] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, *et al.*, "Mastering atari, go, chess and shogi by planning with a learned model," *Nature*, vol. 588, no. 7839, pp. 604–609, 2020.

[76] E. Yilmaz, O. Sanni, M. Kotwicz Herniczek, and B. German, "Deep reinforcement learning approach to air traffic optimization using the muzero algorithm," *AIAA AVIATION 2021 FORUM*, p. 2377, 2021.

[77] A. J. Singh, A. Kumar, and H. C. Lau, "Approximate difference rewards for scalable multigent reinforcement learning," *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 1655–1657, 2021.

[78] R. Isufaj, M. Omeri, and M. A. Piera, "Multi-uav conflict resolution with graph convolutional reinforcement learning," *Applied Sciences*, vol. 12, no. 2, p. 610, 2022.

[79] K. D. Julian, J. Lopez, J. S. Brush, M. P. Owen, and M. J. Kochenderfer, "Policy compression for aircraft collision avoidance systems," *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*, pp. 1–10, 2016.

[80] Y. Xie, A. Gardi, and R. Sabatini, "Reinforcement learning-based flow management techniques for urban air mobility and dense low-altitude air traffic operations," *2021 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC)*, pp. 1–10, 2021.

[81] Y. Tang and Y. Xu, "Multi-agent deep reinforcement learning for solving large-scale air traffic flow management problem: A time-step sequential decision approach," *2021 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC)*, pp. 1–10, 2021.

[82] C. Spatharis, A. Bastas, T. Kravaris, K. Blekas, G. A. Vouros, and J. M. Cordero, "Hierarchical multiagent reinforcement learning schemes for air traffic management," *Neural Computing and Applications*, pp. 1–13, 2021.

[83] T. Duong, K. K. Todi, U. Chaudhary, and H.-L. Truong, "Decentralizing air traffic flow management with blockchain-based reinforcement learning," *Ieee*, pp. 1795–1800, 2019.

[84] A. Estes and M. Ball, "Data-driven planning for ground delay programs," *Transportation Research Record*, vol. 2603, no. 1, pp. 13–20, 2017.

[85] J. Jones, Z. Ellenbogen, and Y. Glina, "Recommending strategic air traffic management initiatives in convective weather," *Fourteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2021), Lexington, MA 02421, USA*, 2021.

[86] M. Bloem and N. Bambos, "Ground delay program analytics with behavioral cloning and inverse reinforcement learning," *Journal of Aerospace Information Systems*, vol. 12, no. 3, pp. 299–313, 2015.

[87] E. George and S. S. Khan, "Reinforcement learning for taxi-out time prediction: An improved q-learning approach," *2015 International Conference on Computing and Network Communications (CoCoNet)*, pp. 757–764, 2015.

[88] A. Gosavii, N. Bandla, and T. K. Das, "A reinforcement learning approach to a single leg airline revenue management problem with multiple fare classes and overbooking," *IIE transactions*, vol. 34, no. 9, pp. 729–742, 2002.

[89] R. J. Lawhead and A. Gosavi, "A bounded actor–critic reinforcement learning algorithm applied to airline revenue management," *Engineering Applications of Artificial Intelligence*, vol. 82, pp. 252–262, 2019.

[90] N. Bondoux, A. Q. Nguyen, T. Fiig, and R. Acuna-Agost, "Reinforcement learning applied to airline revenue management," *Journal of Revenue and Pricing Management*, vol. 19, no. 5, pp. 332–348, 2020.

[91] S. A. Shihab and P. Wei, "A deep reinforcement learning approach to seat inventory control for airline revenue management," *Journal of Revenue and Pricing Management*, pp. 1–17, 2021.

[92] R. Wang, X. Gan, Q. Li, and X. Yan, "Solving a joint pricing and inventory control problem for perishables via deep reinforcement learning," *Complexity*, vol. 2021, 2021.

[93] N. E. Alamdari and G. Savard, "Deep reinforcement learning in seat inventory control problem: an action generation approach," *Journal of Revenue and Pricing Management*, vol. 20, no. 5, pp. 566–579, 2021.

[94] K. T. Talluri, G. Van Ryzin, and G. Van Ryzin, *The theory and practice of revenue management*. Springer, 2004, vol. 1.

[95] P. Belobaba, "Air travel demand and airline seat inventory management," Ph.D. dissertation, Massachusetts Institute of Technology, 1987.

[96] Z. Li, X. Chen, M. Xie, and Z. Zhao, "Adaptive fault-tolerant tracking control of flying-wing unmanned aerial vehicle with system input saturation and state constraints," *Transactions of the Institute of Measurement and Control*, vol. 44, no. 4, pp. 880–891, 2022.

[97] B. Li, Z.-p. Yang, D.-q. Chen, S.-y. Liang, and H. Ma, "Maneuvering target tracking of uav based on mn-ddpg and transfer learning," *Defence Technology*, vol. 17, no. 2, pp. 457–466, 2021.

[98] W. Zhao, H. Chu, X. Miao, L. Guo, H. Shen, C. Zhu, F. Zhang, and D. Liang, "Research on the multiagent joint proximal policy optimization algorithm controlling cooperative fixed-wing uav obstacle avoidance," *Sensors*, vol. 20, no. 16, p. 4546, 2020.

[99] S. Lee, T. Shim, S. Kim, J. Park, K. Hong, and H. Bang, "Vision-based autonomous landing of a multi-copter unmanned aerial vehicle using reinforcement learning," *2018 International Conference on Unmanned Aircraft Systems (ICUAS)*, pp. 108–114, 2018.

[100] B. Xian, X. Zhang, H. Zhang, and X. Gu, "Robust adaptive control for a small unmanned helicopter using reinforcement learning," *IEEE Transactions on Neural Networks and Learning Systems*, 2021.

[101] Y. Zhen, M. Hao, and W. Sun, "Deep reinforcement learning attitude control of fixed-wing uavs," *2020 3rd International Conference on Unmanned Systems (ICUS)*, pp. 239–244, 2020.

[102] D. Huang, J. Hu, Z. Peng, B. Chen, M. Hao, and B. K. Ghosh, "Model-free based reinforcement learning control strategy of aircraft attitude systems," *2020 Chinese Automation Congress (CAC)*, pp. 743–748, 2020.

[103] X. Huang, W. Luo, and J. Liu, "Attitude control of fixed-wing uav based on ddqn," *2019 Chinese Automation Congress (CAC)*, pp. 4722–4726, 2019.

[104] H. Shi, X. Li, K.-S. Hwang, W. Pan, and G. Xu, "Decoupled visual servoing with fuzzy q-learning," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 1, pp. 241–252, 2016.

[105] D. Kim, G. Oh, Y. Seo, and Y. Kim, "Reinforcement learning-based optimal flat spin recovery for unmanned aerial vehicle," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 4, pp. 1076–1084, 2017.

[106] Y. Zuo, K. Deng, Y. Yang, and T. Huang, "Flight attitude simulator control system design based on model-free reinforcement learning method," *2019 IEEE 3rd Advanced Information Management, Communicates, Electronic and Automation Control Conference (IMCEC)*, pp. 355–361, 2019.

[107] Z. Wang, W. Luo, Q. Gong, Y. Cui, R. Tao, Q. Wang, Q. Liang, and S. Wang, "Attitude controller design based on deep reinforcement learning for low-cost aircraft," *2020 Chinese Automation Congress (CAC)*, pp. 463–467, 2020.

[108] E. Bøhn, E. M. Coates, S. Moe, and T. A. Johansen, "Deep reinforcement learning attitude control of fixed-wing uavs using proximal policy optimization," *2019 International Conference on Unmanned Aircraft Systems (ICUAS)*, pp. 523–533, 2019.

[109] D. Wada, S. A. Araujo-Estrada, and S. Windsor, "Unmanned aerial vehicle pitch control under delay using deep reinforcement learning with continuous action in wind tunnel test," *Aerospace*, vol. 8, no. 9, p. 258, 2021.

[110] K. Zhang, K. Li, H. Shi, Z. Zhang, and Z. Liu, "Autonomous guidance maneuver control and decision-making algorithm based on deep reinforcement learning uav route," *Systems Engineering and Electronics*, vol. 42, no. 7, pp. 1567–1574, 2020.

[111] M. Al-Gabalawy, "Machine learning for aircraft control," *Journal of Advanced Research in Dynamical and Control Systems*, vol. 11, pp. 3165–3191, 2019.

[112] W. Koch, R. Mancuso, R. West, and A. Bestavros, "Reinforcement learning for uav attitude control," *ACM Transactions on Cyber-Physical Systems*, vol. 3, no. 2, pp. 1–21, 2019.

[113] Y. Wang, J. Sun, H. He, and C. Sun, "Deterministic policy gradient with integral compensator for robust quadrotor control," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 50, no. 10, pp. 3713–3725, 2019.

[114] F. Fei, Z. Tu, D. Xu, and X. Deng, "Learn-to-recover: Retrofitting uavs with reinforcement learning-assisted flight control under cyber-physical attacks," *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7358–7364, 2020.

[115] L. Bhan, M. Quinones-Grueiro, and G. Biswas, "Fault tolerant control combining reinforcement learning and model-based control," *2021 5th International Conference on Control and Fault-Tolerant Systems (SysTol)*, pp. 31–36, 2021.

[116] M. Blanke, M. Kinnaert, J. Lunze, M. Staroswiecki, and J. Schröder, *Diagnosis and fault-tolerant control*. Springer, 2006, vol. 2.

[117] S. Ragi and E. K. Chong, "Uav path planning in a dynamic environment via partially observable markov decision process," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 4, pp. 2397–2412, 2013.

[118] B. Zhang, Z. Mao, W. Liu, and J. Liu, "Geometric reinforcement learning for path planning of uavs," *Journal of Intelligent & Robotic Systems*, vol. 77, no. 2, pp. 391–409, 2015.

[119] C. Yan, X. Xiang, and C. Wang, "Towards real-time path planning through deep reinforcement learning for a uav in dynamic environments," *Journal of Intelligent & Robotic Systems*, vol. 98, no. 2, pp. 297–309, 2020.

[120] C. Spatharis, T. Kravaris, G. A. Vouros, K. Blekas, G. Chalkiadakis, J. M. C. Garcia, and E. C. Fernandez, "Multiagent reinforcement learning methods to resolve demand capacity balance problems," *Proceedings of the 10th Hellenic Conference on Artificial Intelligence*, pp. 1–9, 2018.

[121] T. Kravaris, C. Spatharis, A. Bastas, G. A. Vouros, K. Blekas, G. Andrienko, N. Andrienko, and J. M. C. Garcia, "Resolving congestions in the air traffic management domain via multiagent reinforcement learning methods," *arXiv preprint arXiv:1912.06860*, 2019.

[122] Y. Chen, Y. Xu, M. Hu, and L. Yang, "Demand and capacity balancing technology based on multi-agent reinforcement learning," *2021 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC)*, pp. 1–9, 2021.

[123] J. Bertram, P. Wei, and J. Zambreno, "Scalable fastmdp for pre-departure airspace reservation and strategic de-conflict," *AIAA Scitech 2021 Forum*, p. 0779, 2021.

[124] A. Wickman, "Exploring feasibility of reinforcement learning flight route planning," Bachelor's Thesis, Department of Computer and Information Science, 2021.

[125] K. Tumer and A. Agogino, "Distributed agent-based air traffic flow management," *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pp. 1–8, 2007.

[126] C. Huang and Y. Xu, "Integrated frameworks of unsupervised, supervised and reinforcement learning for solving air traffic flow management problem," *2021 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC)*, pp. 1–10, 2021.

[127] P. Juntama, D. Delahaye, S. Chaimatanan, and S. Alam, "Hyper-heuristic approach based on reinforcement learning for air traffic complexity mitigation," *Journal of Aerospace Information Systems*, pp. 1–16, 2022.

[128] K. Zhang, Y. Yang, C. Xu, D. Liu, and H. Song, "Learning-to-dispatch: Reinforcement learning based flight planning under emergency," *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pp. 1821–1826, 2021.

[129] Q. Deng and B. F. Santos, "Lookahead approximate dynamic programming for stochastic aircraft maintenance check scheduling optimization," *European Journal of Operational Research*, vol. 299, no. 3, pp. 814–833, 2022.

[130] Y. Hu, X. Miao, J. Zhang, J. Liu, and E. Pan, "Reinforcement learning-driven maintenance strategy: A novel solution for long-term aircraft maintenance decision optimization," *Computers & Industrial Engineering*, vol. 153, p. 107056, 2021.

[131] J. Ribeiro, P. Andrade, M. Carvalho, C. Silva, B. Ribeiro, and L. Roque, "Playful probes for design interaction with machine learning: A tool for aircraft condition-based maintenance planning and visualisation," *Mathematics*, vol. 10, no. 9, p. 1604, 2022.

[132] P. Van Wesel and A. E. Goodloe, "Challenges in the verification of reinforcement learning algorithms," NASA, Tech. Rep., 2017.

[133] A. Baheri, H. Ren, B. Johnson, P. Razzaghi, and P. Wei, "A verification framework for certifying learning-based safety-critical aviation systems," *AIAA AVIATION 2022 Forum*, 2022.

[134] J. Garcıa and F. Fernández, "A comprehensive survey on safe reinforcement learning," *Journal of Machine Learning Research*, vol. 16, no. 1, pp. 1437–1480, 2015.

[135] B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. Al Sallab, S. Yogamani, and P. Pérez, "Deep reinforcement learning for autonomous driving: A survey," *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[136] A. Baheri, S. Nageshrao, H. E. Tseng, I. Kolmanovsky, A. Girard, and D. Filev, "Deep reinforcement learning with enhanced safety for autonomous highway driving," in *2020 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2020, pp. 1550–1555.

[137] A. Baheri, "Safe reinforcement learning with mixture density network, with application to autonomous driving," *Results in Control and Optimization*, vol. 6, p. 100095, 2022.

[138] J. Groot, "Improving safety of vertical manoeuvres in a layered airspace with deep reinforcement learning," Master's thesis, Aerospace Engineering, 2021.

[139] C. Lazarus, J. G. Lopez, and M. J. Kochenderfer, "Runtime safety assurance using reinforcement learning," *2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC)*, pp. 1–9, 2020.

[140] P. Wu and J. Chen, "Comparisons of rrt and mcts for safe assured path planning in urban air mobility," *AIAA SCITECH 2022 Forum*, p. 1841, 2022.

[141] P. Wu, X. Yang, P. Wei, and J. Chen, "Safety assured online guidance with airborne separation for urban air mobility operations in uncertain environments," *IEEE Transactions on Intelligent Transportation Systems*, 2022.

[142] M. Hawley, R. Bharadwaj, and V. Venkataraman, "Real-time mitigation of loss of separation events using reinforcement learning," *2019 IEEE/AIAA 38th Digital Avionics Systems Conference (DASC)*, pp. 1–6, 2019.